

Minutes of X3T11 ad hoc HIPPI, and HNF Technical Committee (TC)
Special Working Meeting to Address HIPPI-6400 Issues Only
Thursday & Friday, May 9 & 10, 1996
Dallas, Texas

1. Opening remarks and introductions

The Chairman, Don Tolmie of Los Alamos National Laboratory, opened this special working meeting at 1 PM by thanking Roger Ronald and E-Systems for hosting this meeting. This group is constituted as both the HIPPI special working group (SWG) under X3T11, and the HIPPI Networking Forum (HNF) - Technical Committee (TC).

2. Review / modify the draft agenda

The draft agenda was distributed by e-mail before the meeting, and hard copies were at the meeting. Additional items were added for:

- Review of previous action items
- Scheduled transfers on VC1 and VC2
- 4b/5b Patent report

3. Review minutes of April 9-10, 1996, Palm Springs, CA, meeting

The HIPPI-6400 April 9-10 minutes were available on the new web page and a few copies were distributed at the meeting. The minutes were reviewed and no additions or corrections were suggested.

Greg Chesson of SGI moved, and John Renwick of NetStar seconded, to approve the minutes as written. Approved unanimously.

4. Review of old action items

The action items from the April, 1996, meeting were reviewed for the current status.

1. Marti Bancroft to draft an annex describing problems and potentials for different VC priority schemes. (Done)
2. Stan Swirhun, and others, to consider problems with FRAME signal frequency. (Carryover)
3. Greg Chesson and others at SGI to consider other codings for the FRAME signal to avoid the long steady periods. (Done)
4. Dave Parry to provide some explanatory text about the 10 μ s requirement for training sequences. (Done)

5. Greg Chesson to check potential CRC codes for goodness, and propose one for the ECRC. (Done)
6. Bob Newhall to provide list of errors that will be detected by SuMAC. (Done - in SuMAC specification, see item 8)
7. Bob Newhall to provide descriptions of what gets reset by Reset and Initialize, and when. (Done - in SuMAC specification, see item 8)
8. Greg Chesson and Jim Pinkerton to merge the MPI and Schedule headers and propose a general micropacket header. (Carryover)
9. Henry Brandt to contact potential copper vendors before the Santa Fe meeting, explain the problem, and solicit their participation in specifying the copper interface. (Done)
10. Don Tolmie to make copper guidelines available via e-mail. (Carryover)
11. Roger Ronald to update HIPPI-6400-SC Rev 0.01 with changes agreed to at this meeting. (Done)
12. Don Tolmie to update HIPPI-6400-PH Rev 0.2 with changes agreed to at this meeting. (Done)

5. Documents on new web page at <http://www.cic-5.lanl.gov/~det/>

Don Tolmie has constructed a HIPPI documents web page as a convenient download site for HIPPI SWG document revisions and other committee documents. Don stated that with the documents available on the web, he plans to bring only a few documents to the meetings, not documents for everyone as done before.

5.1 Is PDF format OK ?

After some discussion about maintaining both PDF and Postscript copies of all documents, it was decided that Don Tolmie would only maintain PDF documents because of the small PDF file size, the ease of creating a PDF, the access to the freeware PDF Acrobat viewer, and the difficulty and work involved in keeping two versions synchronized.

5.2 Other suggestions ?

There were no other suggestions for improving the web pages.

6. Review HIPPI-6400-PH changes since last meeting (Reference Rev 0.21)

HIPPI-6400-PH Rev 0.21 was available via http at the new web page before the meeting, and was also distributed at the meeting. The major changes since Rev 0.20 were listed in the front of the document.

6.1 Changes accepted without changes

- Wording change of "digital data" to "user data" in the Abstract, and the first paragraphs of the Foreword, Introduction, and Scope.
- Last bullet in the Scope.
- Acknowledgment definition.
- Renumbering of first bits and bytes transmitted, and bit assignments for the control bits.
- Control bits summary figure.
- VC3 message size specification of a max of 4 GBytes.
- Deleting the order of VC service.
- Changes to the tables defining the bits and bytes assigned to each signal line.
- Changing the byte numbers in the description of the 4b/5b coding.
- Changes to the figures showing the micropacket formats (namely the byte number changes).
- Text describing the FRAME signal.

6.2 Additional changes were suggested for:

- It is not yet decided that the Admin micropackets will be defined in HIPPI-6400-SC. Still an open issue.
- The contents of the Micropacket contents summary table was modified by changing the heading "Routing Information" to "32 bytes of Header Information (see 6.7)". The TAIL bit for the Null and Credit-only micropackets was changed from 1* to 0*. The VC for an Admin micropacket was

changed from "any" to "Request on VC1, Response on VC2". The notes at the bottom of the table were changed from "...do not check at receiver" to "a receiver must permit any value". The last note was changed from "...ECRC continues from the previous micropacket, it is not initialized" to "...ECRC as defined in 6.6.3".

- Major changes were agreed to for the "Intended use of CRCs", they are not detailed here. The intent is to make more positive statements.

- Various contradictions or confusing statements were discussed that appear throughout clause 6.6. The text describing the ECRC was modified some, and it was noted that the polynomial shown was really the one for the LCRC. The figure showing the implementation example was accepted. Don Tolmie took an action item to revise clause 6.6 and e-mail the revision to the reflector for comments. A place holder for the LCRC stomp-code to be discussed at future meetings was requested. The revisions made to 6.6 will carry over to other parts of the document including 6.1.

- The text describing driving the signal lines to "zero Disparity" at the start of a training sequence was changed to "0 or ± 1 ."

- It was agreed the reversed bit descriptions of Note 4 add confusion to clause 10.1 in its MSB grouping, (dcba). The first two bit groupings will be reversed to show transmission order, (abcd) and (wxTyz).

- The text in 14.1 defining the Souse CLOCK signal was generally accepted, but needs to be expanded for both 16-bit and 8-bit systems.

- The figure in B.1 showing the 4b/5b encoding and decoding will be changed to show more of a serial stream between the Source and Destination.

7. HIPPI-6400-PH Proposals

7.1 Header micropacket contents (6.7, all of 7)

Greg Chesson presented ongoing work at SGI concerning merging the schedule header with the SGI system bypass header. He presented the following parts:

M = "media access control - MAC" portion:
destination address, source address, Ethertype, length
FP = HIPPI-800 framing protocol header
S = schedule header

BP =operating system bypass header (S+X)

He also proposed that a Request to Send (RQS) Clear to Send (CTS) protocol be used to control the VC3 transfers. This terminology was accepted. Proposed scheduled transfer parameters are shown below for the different commands. {Example values are shown in braces}

RQS : (from Originating Source)
[DSIT] (Dest_Port, Source_Port, ID, Tag)
Length {16M} (32-bit)
S_Offset {0} (32-bit)
MessageSize {16k} (32-bit)

RQS_Response : (from Final Destination)
[DSIT] (Dest_Port, Source_Port, ID, Tag)
Length {16M}
D_Offset {0}
MessageSize {16k}Minimum(SRC,DST)

CTS : (from Final Destination)
[DSIT] (Dest_Port, Source_Port, ID, Tag)
Block_Number {0} Striping + Retransmission
Count {1024} of Messages
BUFX {S}

DATA1 : (from Originating Source)
[DSIT] (Dest_Port, Source_Port, ID, Tag)
D_Offset {0+F} Mostly Zero + First Flag
Count {1024} Down to zero
BUFX {S}

DATA2 : (from Originating Source)
[DSIT] (Dest_Port, Source_Port, ID, Tag)
D_Offset {0}
Count {1023}
BUFX {S+1}

Other header issues include:

- RQS + I/FP/D1 + Parameters
- D1 Padding
- unlimited size transfers
- abort by source or destination
- bit values used
- command names
- responses (acks), timeouts, failure and recovery
- camp on

A special working meeting to address just the header issues was called for Wednesday, May 22, at Essential Communications in Albuquerque. Greg felt that they would be far enough along then to propose something more definite.

7.2 LCRC and ECRC simulation results (4.8, 6.6)

Greg Chesson presented his simulation results confirming the strength of the two 16-bit CRC's. There were no cases where both CRC's failed for any errors up to 4-bits. Further simulation and confirmation will stabilize the CRC choice. James Hoffman took an action item to verify the results.

7.3 Retransmission discussion (4.7, 8.2, 8.4, 9.1, 9.2)

The transmission sequencing was discussed during the SuMAC presentation and retransmission was clarified. A short retransmission example is presented below:

At time 0, the Source begins sending micropackets with incrementing TSEQ values. The Destination receives answering RSEQ's confirming reception of the TSEQ up to micropacket 4. At the Destination, micropacket 5 is received in error, and is not acknowledged (time 8). At time 10, the Source's ACK timer for micropacket 5 times out. The Source begins sending all micropackets from the last valid RSEQ. The transmission buffer is cleared up to the highest valid RSEQ value taking into account the circular counting.

TSEQ	0	1	2	3	4	5	6	7	8	9	5	6	7	8	9
RSEQ	FF	FF	FF	0	1	2	3	4	4	4	4	4	4	5	6
time	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14

Although the ACK timer should be tunable on a link basis, 12 microseconds was declared the default value for the standard.

The first sentence of clause 8.4 was deleted. The rest of 8.3 will be updated to reflect the operations in the example above.

7.4 Reset operations (12)

The SuMAC presentation answered many of the questions about reset operations, and the SuMAC document will be used as the basis for the HIPPI-6400-PH document updates.

7.5 Timing (14)

The tolerance for the CLOCK signals was changed from 100 ppm to 200 ppm to reflect the tolerance of the oscillators on the SGI chips. There was no objection to the change. It was noted that the clause should be rewritten to include both 16-bit and 8-bit implementations, it was not currently complete.

7.6 Allowing scheduled traffic on other VC's (6.2)

Greg Chesson asked the committee to allow scheduled traffic to appear on VC1 and VC2. The group concluded that VC3 requires scheduled traffic and VC1 and VC2 should optionally allow scheduled traffic. The scheduled traffic handshake (always on VC0) must now include a VC identifier.

8. SuMAC chip discussion

Jim Davis of E-Systems presented the SGI SuMAC (Super HIPPI Media Access Controller) preliminary specification. The topics covered are listed below. The specification was copyrighted, hence cannot be distributed by the committee. For a copy of the specification, please contact Dave Parry of SGI at parry@hobbs.asd.sgi.com.

Architecture Overview

Operational Modes

"System" Design Topics

- Warm Reset-propagating reset
- Link Reset
- Credit Control
- Link Retry
- Message Handling
- Crossbar Routing
- Deadlock Prevention
- Link Shutdown
- Buffer Sizing

External Interfaces

The presentation answered many questions concerning: micropacket sequencing and retransmission; message handling, Admin block; overall HIPPI-6400 architecture; resets; timing; and error handling. The next revision of HIPPI-6400-PH will try to incorporate non-implementation-specific details from the SuMAC specification.

It was noted that the SuMAC incorporates an interesting VC crediting scheme whose intention should be incorporated into the document.

9. HIPPI-6400-SC changes since last meeting **(Reference Rev 0.10)**

The following revisions were applied to the switch control document

9.1 "Optional" error checking (5.4)

The group decided that labeling clause 5.4 as optional, might misdirect hardware implementors

regarding error checking at the destination. Possibly the clause should be broken into required and optional error checking.

9.2 Micropacket interleaving (5.5.1)

Micropacket interleaving granularity was discussed and Roger Ronald took an action item to draft an annex describing the advantages and disadvantages of the two extremes, i.e., every micropacket versus every 65 micropackets.

10. HP Patent

Francois Gaulier spoke with patent Hewlett-Packard patent representative Pat Petrino regarding the similarity between HIPPI-6400's 4b/5b encoding scheme to HP's G-link encoding used in HIPPI-Serial. Pat Petrino will contact Don Tolmie to confirm the lack of patent infringement.

11. HIPPI-6400-SC proposals

11.1 Alternate Pathing

Roger Ronald presented his views on a need for alternate pathing through the switching network. Alternate pathing will function on a per message basis allowing out of order message delivery. Alternate pathing requires switch manipulation, hence no alternate pathing is the initial switch setup. Alternate pathing allows multiple switch source ports for a given destination port which may also lend itself to multicast. Roger Ronald took an action item to present an alternate pathing proposal via e-mail.

11.2 Routing and loopback (5.1, 5.2)

The address lookup routing methods were discussed to verify a loopback switching capability. It was determined that each Destination port on a switch has its own 16-bit lookup table allowing loopback.

11.3 Admin micropacket functions

The admin micropacket functions were discussed and Roger Ronald took an action item to organize the required functions for switch control. Items that should be addressed regarding admin micropackets:

- self discovery (auto configuration)
- addressing
- format
- neighbor discovery (bootstrapping)

It was decided that the basic format of the admin micropacket belongs in the HIPPI-6400-PH specification, but details would be provided by the HIPPI-6400-SC.

11.4 SD-3 Project Proposal

Don Tolmie presented the HIPPI-6400-SC SD-3 for forwarding at the X3T11 Santa Fe meeting. No revisions were requested.

12. Status report on copper

Henry Brandt of IBM reported on his communications to cable and connector vendors. He invited multiple companies to attend the upcoming Santa Fe meeting. He also investigated related standards activities such as SCI, FC0 and Gigabit Ethernet.

One interesting result from his investigation is a gigabaud single-bit, duplex I/O signaling chip that only cost \$20 each from Technitrol. Technitrol will be represented in Santa Fe.

13. Other open issues

The HIPPI-6400-PH document was reviewed to see if some of the "Open Issues" could be closed. Many of them were, and will be removed from future revisions. Some, referenced by their clause number, that resulted in discussion included:

3.1.2 - The definition for "BER" will be removed as not being applicable for this standard.

3.1.x - We need to add definitions for "Originating Source" and "Final Destination", and include a figure showing HIPPI-6400 links and switches.

Figure 1 - Will be changed from showing the data lines as "n" to "16(8)". The control lines will be changed in a like manner.

6.3 - The question was raised about handling unknown TYPE's. To preserve future growth and keep current convention, TYPE's 0-7 will be dropped at the link layer except for the RSEQ field which will be used. TYPE's 8-15 will be forwarded to upper layers.

6.7 - It was agreed to require all unused data bytes at the end of a message to be transmitted as zeroes for security reasons.

14. Future meeting schedule

May 22, 1996, Albuquerque, NM

A meeting specifically to discuss the HIPPI-6400 Header format and implications including OS-Bypass and legacy HIPPI mapping will be held in Albuquerque, NM, 1 pm to 9 pm. Michael McGowen and Essential Communications will be the host.

June 10-11, 1996, Santa Fe, NM

Monday, June 10 -

9 AM - 9 PM -- HIPPI-6400

Tuesday, June 11 -

8 AM - 3 PM -- HIPPI General, HNF, HIPPI-6400

3 PM - 6 PM -- HIPPI-6400 Optical

6 PM - 9 PM -- HIPPI-6400 Copper

The location is the La Fonda Hotel, 100 E. San Francisco, Santa Fe, NM 87501, phone (505) 982-5511, Fax (505) 988-2952. The rate is \$105 single or double, parking and tax included. The closing data for reservations is May 10. The group name is X3T11-Los Alamos. Don Tolmie and Los Alamos are the host.

The schedule for the rest of 1996 is listed below. The Plenary meetings include HIPPI-6400, an HNF plenary, and all other HIPPI items. The Interim meeting cover just HIPPI-6400 items.

Jul 10-11	Interim	Mountain View, CA	SGI
Aug 5-6	Plenary	Honolulu, HI	Hitachi
Sep 11-12	Interim	Albuquerque, NM	Los Alamos
Oct 7-8	Plenary	St. Petersburg Beach, FL	AMP
Nov 6-7	Interim	Phoenix, AZ	Loral
Dec 2-3	Plenary	Rochester, MN	IBM

15. Review action items

1. Stan Swirhun, and others, to consider problems with FRAME signal frequency in optical implementations.
2. Greg Chesson and Jim Pinkerton to merge the MPI and Schedule headers and propose a general micropacket header.
3. Don Tolmie to make copper guidelines available via e-mail.
4. Don Tolmie to submit a rewrite of section 6.6 via e-mail.
5. Greg Chesson to provide Don Tolmie with the parallel CRC equations to be added as an annex.
6. James Hoffman to verify CRC error protection results.

HIPPI-6400 SWG meeting minutes - May 9-10, 1996

7. Greg Chesson to gather SGI header ideas to be used at the May 22 meeting.
8. Roger Ronald to draft an annex on micropacket interleaving to be reviewed via e-mail.
9. Roger Ronald to present a proposal for alternate pathing to e-mail.
10. Roger Ronald to define admin micropacket contents and requirements and present to e-mail.
11. Henry Brandt to inform and invite various optical and copper companies to the Santa Fe meeting.
12. Roger Ronald to update HIPPI-6400-SC Rev 0.10 with changes agreed to at this meeting.
13. Don Tolmie to update HIPPI-6400-PH Rev 0.21 with changes agreed to at this meeting and incorporating details from the SuMAC specification.
14. Greg Chesson to put out a meeting notice for the July meeting at SGI via e-mail.
15. Michael McGowen to put out a meeting notice for the May 22 meeting at Essential via e-mail.

Attendance

Marti Bancroft	Cray Research/SGI
Mark Kelley	Cray Research/SGI
Bill Kranda	Essential Communications
Michael McGowen	Essential Communications
John Armstrong	E-Systems
Alex Brown	E-Systems
Craig Davidson	E-Systems
Jim Davis	E-Systems
Mike Ficarra	E-Systems
Sam Locke	E-Systems
Mike Montana	E-Systems
Roger Ronald	E-Systems
Francois Gaullier	Hewlett-Packard
Henry Brandt	IBM
Chris Olson	Lockheed-Martin
James Hoffman	Los Alamos National Lab
Wally St. John	Los Alamos National Lab
Don Tolmie	Los Alamos National Lab
Stephen Quan	NASA Ames Research
Von Welch	NCSA
John Renwick	NetStar
Joe Parker	Optivision
Greg Chesson	Silicon Graphics
Larry McVoy	Silicon Graphics
Bob Newhall	Silicon Graphics

16. Adjournment

The meeting adjourned at 2 PM on the second day after a well organized, productive meeting.